4-2013

A Cognitive Approach to Vision for a Mobile Robot

Paul Benjamin

Pace University

Christopher Funk

Pace University

Damian M. Lyons

Fordham University

Follow this and additional works at: https://fordham.bepress.com/frcv_facultypubs

Part of the [Robotics Commons](#)

Recommended Citation

Benjamin, Paul; Funk, Christopher; and Lyons, Damian M., "A Cognitive Approach to Vision for a Mobile Robot" (2013). *Faculty Publications*. 34.

https://fordham.bepress.com/frcv_facultypubs/34

This Article is brought to you for free and open access by the Robotics and Computer Vision Laboratory at DigitalResearch@Fordham. It has been accepted for inclusion in Faculty Publications by an authorized administrator of DigitalResearch@Fordham. For more information, please contact considine@fordham.edu.

A Cognitive Approach to Vision for a Mobile Robot

D. Paul Benjamin
Christopher Funk

Pace University, 1 Pace Plaza, New York, New York 10038, 212-346-1012
benjamin@pace.edu

Damian Lyons

Fordham University, 340 JMH, 441 E. Fordham Rd., Bronx, NY 10458, 718-817-4485
dlyons@fordham.edu

ABSTRACT

We describe a cognitive vision system for a mobile robot. This system works in a manner similar to the human vision system, using saccadic, vergence and pursuit movements to extract information from visual input. At each fixation, the system builds a 3D model of a small region, combining information about distance, shape, texture and motion. These 3D models are embedded within an overall 3D model of the robot's environment. This approach turns the computer vision problem into a search problem, with the goal of constructing a physically realistic model of the entire environment.

At each step, the vision system selects a point in the visual input to focus on. The distance, shape, texture and motion information are computed in a small region and used to build a mesh in a 3D virtual world. Background knowledge is used to extend this structure as appropriate, e.g. if a patch of wall is seen, it is hypothesized to be part of a large wall and the entire wall is created in the virtual world, or if part of an object is recognized, the whole object's mesh is retrieved from the library of objects and placed into the virtual world. The difference between the input from the real camera and from the virtual camera is compared using local Gaussians, creating an error mask that indicates the main differences between them. This is then used to select the next points to focus on.

This approach permits us to use very expensive algorithms on small localities, thus generating very accurate models. It also is task-oriented, permitting the robot to use its knowledge about its task and goals to decide which parts of the environment need to be examined.

The software components of this architecture include PhysX for the 3D virtual world, OpenCV and the Point Cloud Library for visual processing, and the Soar cognitive architecture, which controls the perceptual processing and robot planning. The hardware is a custom-built pan-tilt stereo color camera.

We describe experiments using both static and moving objects.

Keywords: robot schemas, virtual world, Soar cognitive architecture, port automata

1. INTRODUCTION

The current generation of behavior-based robots is programmed directly for each task. The programs are written in a way that uses as few built-in cognitive assumptions as possible, and as much sensory information as possible. The lack of cognitive assumptions gives them a certain robustness and generality in dealing with unstructured environments. However it is proving a challenge to extend the competence of such systems beyond navigation and some simple tasks [11]. Complex tasks that involve reasoning about spatial and temporal relationships require robots to possess more advanced mechanisms for planning, reasoning, learning and representation.

The ADAPT project (Adaptive Dynamics and Active Perception for Thought) is a collaboration of three university research

groups at Pace University, Brigham Young University, and Fordham University to produce a robot cognitive architecture that integrates the structures designed by cognitive scientists with those developed by robotics researchers for real-time perception and control [1-4]. Our goal is to create a new kind of robot architecture capable of robust behavior in unstructured environments, exhibiting problem solving and planning skills, learning from experience, novel methods of perception, comprehension of natural language and speech generation.

Our approach is fundamentally different from other hybrid architectures, which typically attempt to build a comprehensive system by connecting modules for each different capability: learning, vision, natural language, etc. Instead, we are building a *complete cognitive robotic architecture* by merging RS [9,10], which provides a model for building and reasoning about sensory-motor schemas, with Soar [5,6], a cognitive architecture that is under development at a number of universities. RS possesses a sophisticated formal language for reasoning about networks of port automata and has been successfully applied to robot planning [8]. Soar is a unified cognitive architecture [45] that has been successfully applied to a wide range of tasks including tactical air warfare [12].

One of the most unique and important aspects of our architecture is its treatment of perception and language and their relationship to knowledge representation. Our view of perception is that it is an "active" process [7] that is goal-directed and task-dependent, i.e. it is a cognitive problem-solving process rather than a peripheral activity separate from cognitive processing. Furthermore, we view perception as intimately linked with the formation and modification of symbolic representations, so that perception's main purpose is to identify, build and modify representational structures that make the robot's goals easier to achieve. In this view, the robot solves problems primarily by searching among different ways of perceiving the world and the task. This is in contrast to the usual approach of searching among sequences of actions in one or a few fixed representations. Each way of perceiving the world and task leads to a distinct symbolic formulation.

This cognitive approach to vision casts visual perception as a search process. ADAPT searches among ways to build a model of itself and the environment; its goal is a model that is sufficiently accurate to use for planning. The next section describes this search process in detail.

2. A COGNITIVE APPROACH TO VISION

One of the distinctive features of ADAPT is its virtual world. ADAPT's virtual world is a multimedia simulation platform capable of realistic simulations of physical phenomena. It combines the various forms of map information found in most robots: topological, metric and conceptual information. In addition, this virtual world has a sophisticated physics plugin, giving it the ability to predict dynamics. ADAPT completely controls this virtual world, and can create arbitrary objects and behaviors in it. Central to ADAPT's use of its virtual world is its ability to view these constructions from any point.

ADAPT uses this virtual world in a novel way. Typical robotics architectures connect their sensory mechanisms to their world models, so that sensory data is processed and modeled in the world model. The reasoning engine then operates on the world model to plan the robot's behaviors. This type of architecture treats perception as a separate process from the central reasoning, and typically the implementation reflects this, e.g. a computer vision module processes the vision data and puts symbolic representations of the recognized objects and their relationships in the world model, and the reasoning engine then manipulates these symbols to plan and learn. The reasoning engine does not process the sensory data.

In contrast, ADAPT's virtual world is not connected to its sensory processes. ADAPT's sensory data is placed directly in the reasoning engine (after some low-level processing); the reasoning engine's principal task in ADAPT is to reason about how to model the data. It does this using the following loop:

It compares visual data from the camera with visual data from the corresponding virtual camera, using a least-squares measure to find areas of disagreement. Each disagreement causes a Soar operator to be proposed to attend to that disagreement.

One Soar operator is selected, based on the robot's current goals. For example, if the goal is navigation, operators will be preferred that are for visual disagreements in the robot's current path. The selected operator fires, causing the cameras to saccade to the area of disagreement and fixate on it.

Stereo disparity, color segmentation, and optical flow are computed only in the small region of focus. Restricting the computation to this small region permits the use of highly accurate but computationally expensive algorithms for these computations, e.g. disparity of disparities.

This information is input to the object recognition database, and a mesh model of the best match is rendered into the virtual world. If the information indicates that a current mesh model is inaccurate, that model is modified to incorporate the new information. Optical flow is used to create or update the RS process model associated with each object.

The reasoning engine searches alternative combinations of virtual entities and behaviors to attempt to minimize the measured disagreement. In this way, *perception becomes a problem-solving process*. This enables all the knowledge of the system to be brought to bear on perception, and unifies the reasoning and learning processes of problem solving with those of perception.

The research hypothesis of this work is that the movements of the robot's cameras are only those that are necessary to build a sufficiently accurate world model for the robot's current goals. For example, if the goal is to navigate through a room, the model needs to contain any obstacles that would be encountered, giving their approximate positions and sizes. Other information does not need to be rendered into the virtual world, so this approach trades model accuracy for speed.

In this way, ADAPT prunes the information at the perception stage, using its knowledge about objects and dynamics. This is in contrast to the usual approach gathering lots of sensory information, processing it all and rendering it into a world model, then deciding which information is necessary for decision making. This latter approach wastes a great deal of time and processing effort to refine information that is discarded in the decision making process.

3. EXAMPLE

A robot wishes to predict the path of a bouncing ball so that it can intercept the ball as quickly as possible. The ball will bounce off walls that will alter its path. The robot needs to perceive the objects that the ball will hit and also perceive the ball's motion, then combine this information to produce a predicted path.

In Figure 1 below, we see two boards placed on the floor. Our vision system detects keypoints and lines in the images, then selects a region for initial focus. The density of keypoints at the bottom of the images causes this to be selected as the region of focus. This region is denoted by the dark boxes in Figure 2.

Figure 1. Initial processing of two boards on the floor, showing keypoints and lines.

The upper corners of the boards are detected, and registration of the left and right images produces an initial correspondence, together with position and orientation data.

Figure 2. The region of focus is the boxes at the lower left, yielding the correspondence between the corners.

A saccade is performed to the next region of focus. In this case, the closest region is chosen, which is a correspondence between the right boards. This process is repeated four times until the top corners of the boards are reached.

Segmentation information is added, producing additional correspondences. This is combined with the correspondences from the keypoints, yielding a small set of best correspondences. In Figure 3, we see the tops of the boards correspond.

Figure 3. Segmentation correspondence between the boards.

Finally, the boards are rendered in PhysX, and a ball is added. The ball is rolled from right to left.

Figure 4. A ball is rolled between two boards. Left and right images are at top. The virtual world is at bottom.

The direction and velocity of the ball are computed over a small interval then duplicated in the virtual world. The physics engine is then run much faster than real-time, producing a predicted path for the ball.

A number of videos showing this process in various scenarios are available at <http://csis.pace.edu/robotlab/videos.html>

Figure 5. Schematic of the information flow in the system.

4. CURRENT WORK

We are enhancing the interaction between the virtual world and the camera input. At present, they are compared using the MMD (Match-Mediated Difference) which uses a local Gaussian to find the largest differences between the image from the real camera and the image from a virtual camera situated at the same location in the virtual world. We are moving toward extracting higher-level features from the virtual world to be compared to the camera input. In particular, we are implementing a hierarchical object detection system that will traverse a hierarchy of features found in the virtual world and try to detect them in the real image.

In addition, we are now beginning the design of skeleton detection, and the rendering of realistic human figures.

5. SUMMARY

We have sketched the overall design of a cognitive computer vision system based on the structure and behavior of the human visual system. Our system builds a 3D model of a dynamic environment, updating it in real time as the world changes. Stereo cameras are moved and refocused by a cognitive architecture to build and update this model. We are investigating the accuracy and efficiency of different strategies for building this model, and evaluating numerous tradeoffs in designing such a system. In addition, we are exploring the cognitive plausibility this approach by comparing the performance and behavior with human saccade data. The vision system will be used on a mobile robot that performs a variety of tasks and interacts with humans.

Further information on this work, including video clips showing the robot moving under the control of schemas and the use of the world model, can be downloaded from the website for the Pace University Robotics Lab: <http://esis.pace.edu/robotlab>

REFERENCES

- [1] Benjamin, D. Paul, Damian Lyons and Deryle Lonsdale, "Embodying a Cognitive Model in a Mobile Robot", Proceedings of the SPIE Conference on Intelligent Robots and Computer Vision, Boston, October, 2006.
- [2] Benjamin, D. Paul, Damian Lyons and Thomas Achtemichuk, "Obstacle Avoidance using Predictive Vision based on a Dynamic 3D World Model", Proceedings of the SPIE Conference on Intelligent Robots and Computer Vision, Boston, October, 2006.
- [3] Benjamin, D. Paul, Damian Lyons and Deryle Lonsdale, "Designing a Robot Cognitive Architecture with Concurrency and Active Perception", Proceedings of the AAAI Fall Symposium on the Intersection of Cognitive Science and Robotics, Washington, D.C., October, 2004.
- [4] Benjamin, D. Paul, Damian Lyons and Deryle Lonsdale, "Cognitive Robots: Integrating Perception, Action and Problem Solving in Behavior-Based Robots", AAMAS-2004 Proceedings, pp. 1308-1309, 2004.
- [5] Laird, J.E., Newell, A. and Rosenbloom, P.S., "Soar: An Architecture for General Intelligence", *Artificial Intelligence* **33**, pp.1-64, 1987.
- [6] Newell, Allen, *Unified Theories of Cognition*, Harvard University Press, Cambridge, Massachusetts, 1990.
- [7] Blake, A. and Yuille, A., eds, *Active Vision*, MIT Press, Cambridge, MA, 1992.
- [8] Lyons, D.M. and Hendriks, A., "Exploiting Patterns of Interaction to Select Reactions", Special Issue on Computational Theories of Interaction, *Artificial Intelligence* **73**, 1995, pp.117-148.
- [9] Lyons, D.M., "Representing and Analysing Action Plans as Networks of Concurrent Processes", *IEEE Transactions on Robotics and Automation*, June 1993.

- [10] Lyons, D.M. and Arbib, M.A., "A Formal Model of Computation for Sensory-based Robotics", IEEE Transactions on Robotics and Automation **5**(3), Jun. 1989.
- [11] Niclescu, M. and Mataric, M., "Extending Behavior-based System Capabilities Using an Abstract Behavior Representation", *Working Notes of the AAAI Fall Symposium on Parallel Cognition*, pages 27-34, North Falmouth, MA, November 3-5, 2000.
- [12] Rosenbloom, P.S., Johnson, W.L., Jones, R.M., Koss, F., Laird, J.E., Lehman, J.F., Rubinoff, R., Schwamb, K.B., and Tambe, M., "Intelligent Automated Agents for Tactical Air Simulation: A Progress Report", Proceedings of the Fourth Conference on Computer Generated Forces and Behavioral Representation, pp.69-78, 1994.